Privacy & Security Workgroup
Draft Transcript
April 1, 2010

Presentation

Judy Sparrow – Office of the National Coordinator – Executive Director
Thank you. Good morning, everybody, and welcome to the Privacy & Security Workgroup call. This is a federal advisory committee, so there will be opportunity at the close of the meeting for the public to make comments. And also, this call will feature a presentation, a Webinar presentation by Oasis. But first, let me do a roll call of the privacy and security workgroup members. Dixie Baker?

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Steve Findlay?

Steve Findlay – Consumers Union – Senior Healthcare Policy Analyst
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Anne Castro?

Anne Castro – BlueCross BlueShield South Carolina – Chief Design Architect
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Dave McCallie?

David McCallie – Cerner Corporation – Vice President of Medical Informatics
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Gina Perez? Wes Rishel? Sharon Terri? Walter Suarez?

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Aneesh Chopra? Jodi Daniel? Joy Pritts? John Moehrke?

John Moehrke – Interoperability & Security, GE – Principal Engineer
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Kevin Stein?

Kevin Stein
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
Joyce DuBow?

Joyce DuBow – AARP Public Policy Institute – Associate Director
Here.

Judy Sparrow – Office of the National Coordinator – Executive Director
And we did invite members from the policy committee's privacy and security workgroup. I believe Deven McGraw is on the phone.

Deven McGraw - Center for Democracy & Technology – Director
Yes.

Judy Sparrow – Office of the National Coordinator – Executive Director
Anyone else from the policy committee workgroup?

Deven McGraw - Center for Democracy & Technology – Director
Joyce is on it too.

Judy Sparrow – Office of the National Coordinator – Executive Director
Yes.

Deven McGraw - Center for Democracy & Technology – Director
She's pulling double duty.

Judy Sparrow – Office of the National Coordinator – Executive Director
Great. Okay. And I know, Dixie, we have Mr. Sabo and Mr. Willett from OASIS, so I'll turn it over to you.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
All right. I'd like to welcome you guys and thank you for dialing in. This is the first of a series of at least four, perhaps more, educational sessions that are focusing on various standardization activities around the topic of consent management, in other words, the management of patient permissions that are given with respect to their personal health information.

I thought that for this first session, the permissions management reference model that's being developed within the OASIS, which is the Organization for the Advancement of Structured Information Standards, would be a good kind of framework for the sessions that follow. So I've invited John Sabo, Michael Willett, and Dee Schur to present that work to us today. I'll let each of you. I know you're from different organizations, so I'll let each of you introduce yourself. John?

John Sabo – OASIS
Dixie, thank you very much, and thanks to the committee and the working group for giving us the opportunity to provide this overview of new work being introduced into OASIS, which I do think will be of interest to the workgroup and to the committee. Probably novel in that will sort of present a new perspective on privacy management, permissions management consent in the privacy context. I work for CA, formally Computer Associates. Michael Willett with WillettWorks, he and I have been longstanding members of the International Security Trust & Privacy Alliance, which is a nonprofit that has been

developing this reference model for a number of years. We are going to split the presentation into two parts, and I'll start with introduction overviews, that type of thing, and then Michael will get into the meats of the presentation.

I am a member – I have longstanding experience in data privacy beginning with Social Security where I worked for 23 years, and also my industry work, and I'm currently a member of the homeland security, data privacy and integrity advisory committee, which focuses mostly on VHS systems from a privacy and data integrity perspective. I'll let Michael introduce his own background briefly. He and I were colleagues at IBM together and, over the last several years, working in ISTPA to develop this model, which is now being introduced into OASIS as part of a standards development process.

The key objectives for us today are to introduce the model. That is, it's a new approach to understanding privacy and it doesn't get into the policies as much as it gets into how to manage and deliver the policies, so it's an unusual approach and we're going to try to introduce it with enough detail in the short time that we have to generate some interest. We'd also like to generate interest in members of the workgroup committees and your companies and industries and in government to participate with OASIS and become part of this new technical committee, which will develop the model.

Our focus will be on use case development, improving the model, and making it more workable in particular sectors, industries, such as health IT, such as smart grid, such as cloud computing. We formed a discussion group that's on the slide here where you can join if there is discussion. Our current intent is to actually initiate the technical committee in May, so we're moving now to establish this, and we'll get into more details later.

On the download section of the Webinar page, you can download this, obviously this presentation, and you can also download the two documents that we will be discussing: the ISTPA privacy management reference model, which will be introduced into OASIS, and another document, the analysis of privacy principles we developed and published in 2007. I think those would be of interest for those on the seminar today.

OASIS itself, just a very quick overview, OASIS is an international standards organization, but not in the sense of the ISO, which relies on national bodies, but as government industry and expert standards body, a nonprofit consortium, which develops and adopts open standards, and it has more than 6,000 participants and 100 organizations in over 100 countries. It has a very clear, transparent, and structured governance process with the standards themselves, although the process for development is lightweight in the sense that it's more expeditious and flexible than some of the international processes. Nevertheless, it allows the development of some very important standards, a number of which have now been adopted by ISO and other international bodies, so I think the OASIS organization is very valuable.

The standards development organization, particularly for standards that rely on XML and a number of very key security standards, such as security assertion markup language, so it's a very fast track, but also very rigid and structured process for the actual implementation and voting and development of standards. There's information on the slide deck here. Dee Schur from OASIS is on the call, and if there are any questions later from the workgroup or committee members about OASIS, she can answer those questions.

Why do we need a privacy management reference model? Why are we talking to you about this today, and what value would it have to you? What I'm going to touch on is not – I'm not a healthcare expert. I'm a privacy expert. I'm not a health IT specialist. Many, clearly you are, and I'm not going to attempt to talk about your challenges with respect to managing privacy. But in the context of health IT, you'll see that

the issues around privacy managements span multiple vertical sectors of the economy and multiple infrastructures, and it's becoming an increasingly large challenge, so we do need a structured model, we believe, in order to begin managing privacy and not just talking about policy statements around privacy.

When we began our work a number of years ago, you looked at the overall privacy landscape, and this slide pretty much, this tiny fraction of the privacy laws, rules, directives, and international compact, whether it's the U.S. Privacy Act of 1974 or really the first moving on to OECD and moving on to various, the … data protection directive, moving on to Safe Harbor national laws … Canada, etc., California Bill 1386 dealing with security breach notification, a state sponsored law, which has had huge ramifications in the U.S., etc. We have things like the AICPA/CICA privacy framework now called the generally accepted privacy practices and principles, which is a very good framework for listing control statements associated with privacy management. So this is a complex set of requirement statements often written in legislative and regulatory language.

But from an implementation or operational perspective, these principles and practices are similar, but they're not standardized. So if you look at the OECD guidelines list of principles, practices against Australia, against APAC, and we could have looked at dozens of these, you'll see some words that are equivalent, some similar. And when you read the definitions that are contained in these instruments internationally, you'll see huge variations in vocabulary, in definition, and in meaning. And you'll see a huge richness in some, and a very much more simple definitions in others, particularly the older instruments.

We, and the little arrow in the right box just indicates there are many, many more of these, but there is the ability to look for commonality, and in the document that you can download, I referenced analysis and privacy principles, an operational study. We looked across 12 of these international instruments, and using researchers and our own expert contributors, we basically looked for a common set, not an attempt to define this for the privacy community, but to look at it from the perspective of an implementer. In other words, if you're wanting, building an architecture in your company or in a sector, if you're looking at health information exchanges, if you're looking at moving data from providers to consumers or to custodians of data in a very complex environment, are there common terms and common understandings from an operational perspective that help guide an architecture.

We can up with this set of accountability notice, consent, etc. You see the list here. We found anonymity contained in a number of the international instruments, data flow across policy boundaries, across system boundaries as a significant component of a number of these, and obviously sensitivity, certain information considered more sensitive and needing additional controls than typical PI or PII, so this document is downloadable. And it's not so much that it is a definitive finding with respect to this set of principles, practices. It's that it offers an interesting methodology to help define those. That might be a useful methodology within the health IT community, as you look at EHRs, PHRs, and all the complex infrastructures that are now being built.

Another key issue is security. We have made the distinction in this work from the beginning that security is critical to privacy. Clearly security, as you saw in those standards listings that I had shown you earlier, and I bolded those, you'll see security is essentially contained in all of them. It's critical—it looks like this thing is moving a little too quickly—security is critical to privacy, and yet it is a component of privacy and distinct. And we see that there are fundamental security services, confidentiality, data integrity, availability. These are typically seen in the community as the key fundamental, foundational components of information security. And we have a lot of standards to implement security in the security discipline.

It's a professional discipline. We have encryption standards, such as the NIST advanced encryption standard. We have security assertion markup language. We have industry standards used in the financial community, payment card industry, data security standard, and we have international standards, the ISO series and dozens and dozens of security standards that are relevant. And a number of them are in and have been developed and are implemented day-to-day that are based on OASIS standards and based on other standards.

We have a very rich discipline. We have a very mature profession. It may not be perfect, and we may have a lot of issues around managing security and risk with security mechanisms. But, nevertheless, we have those controls. We have those mechanisms. We have products, and we have solutions.

The other thing that we have are key security mechanisms, and this is the level, which is above solutions, which is above products, but these are mechanisms that support privacy, and there are many of them. This is especially important because it directly relates to data privacy in terms of individual access and access controls. It's identity lifecycle management and compliance. That is the right people having access to the correct information in a well-defined identity system.

Web access management, federated security mechanisms, service oriented architecture security, this enables trust among multiple entities, so I know in the HIT arena and infrastructures certainly trust and this federated idea of insuring entities have appropriate controls and rights of access is important. Resource protection, you know, controlling and monitoring the use of information by privileged users. You may have the right to go into a system, but you may not have the right to see certain data. We have those technologies. We have data protection technologies for protecting it at rest. We have encryption, and we have the ability to audit, to manage what's going on and provide some audit controls.

The infrastructure is also available. Again, this is all related to security. If you look at this compliance infrastructure, and without getting into great detail, you'll see that we have the ability to take the access management controls and the user management and all of these things and basically integrate them into a set of common policies, workflows, reporting, and so on, and to build these into both your enterprise infrastructure at a facility or at a hospital, as well as a federated infrastructure that crossed all those.

Then on top of—embedded in that, of course, would be the actual applications the various platform services, and the information itself, so we have a compliance and security management infrastructure that's very mature and that is used today, and that can be applied in the integrated health systems that are coming to us. What we don't have, however, is the equivalent for privacy, and the privacy drivers and the privacy issues, of which security is a component, need a similar overlay. They need a similar compliance and management infrastructure, not just compliance, but also management.

The drivers are clear. I mean, you're aware of this in the work you're doing as part of the committee, and that is, you've got networks. We have a very extendible, if not infinite lifecycle of personal and personally identifiable information. It's virtually boundless. PI is not just stored at one entity. It's stored in multiple sites across multiple storage systems. We have multiple principles, legislation, and policies, just as I have explained earlier.

We have the need for operational privacy management standards, not security standards. We have many of them, and the two need to be linked. But from a privacy perspective, we need those, and this is a driver.

Finally, what was a huge motivator the last few years in developing this model, as we move it into OASIS, is the recognition that the new business models and the new mandates for infrastructures, whether it's

social networking systems, electronic government, cloud computing, which is turning out to be a huge, huge economic driver, as well as a technical issue, smart grid, and health IT systems. These new infrastructures will require a similar, overarching management and compliance model. And it will need to make use of various mechanisms and various technologies, and various solutions and products, as well as the appropriate policies and business processes. But these models are now exposing some of the flaws we've had in terms of our ability to manage privacy across policy boundaries, as well as – let's see if I can – one second here … advance the slide.

These challenges, maybe just quickly just show you a few illustrations of what we're talking about. Cloud computing, CA as a company, were involved in the World Economic Forum. I was able to get permission to use this. They're doing a major study on cloud computing deployment. They've held workshops globally in China, India, Europe, and so on. Basically, the see huge economic benefits from an entrepreneurial perspective, but also from the developing countries in this move towards cloud computing. It's going to be a leapfrog in technology if it's properly implemented and well adopted.

But they also found in all these, in the data surveys they've done and workshops, huge perceptions of barriers to the uptake of cloud computing, and there's a major report that's going to be issued in the spring at the World Economic Forum, and some additional work going on. Privacy was the top barrier; 63% of the respondents across all these surveys. Data governance, data ownership, data transfer, cross boarder, and security, so the top issues related to being barriers to the deployment of cloud computing, which everyone expects to have huge benefits, are the very barriers that we've been trying to address through, particularly with privacy in this reference model work.

If you look at smart grid, another example where some of you may be aware of the work NIST is doing in conjunction with smart grid to identify security and privacy risks and the controls that are needed to address those risks, and without getting into the details of this, you've got a huge sense from many people in the privacy community, as well as those who are trying to evaluate risk of integrating devices at the consumer level in the homes, small business into a very complex grid, and there's a link there to the smart grid work that NIST is conducting right now, including the ability to look at some of their privacy findings. Typically, in the privacy area, they have generalized risk findings.

If you look at the security aspects of what they're doing, as I kind of alluded to earlier, huge reams of information about security risk and security controls and, frankly, the privacy side talks about risks generally, but doesn't have that richness of detail because we just don't have that infrastructure in place. We don't have those standards. We don't have those mechanisms.

If you look at the NIST smart grid conceptual model for the grid, and you look at the, and this is out of their work, which you can download, it's attempting to show the relationship among operations and the service providers and the consumers. The far right lower box shows the premises networks. A huge amount of data will be collected, and the aggregated data and individual data will include sensitive PII. And, therefore, there's a desire and need to begin managing that. Again, we don't have the model.

Finally, in your area, network health IT, I'll just use this slide as an expression of the probably minimal complexity. Because it's just showing the key participants from … perspective in the health IT network health IT system, but again it's the same question. You're moving patient information, patient data, sensitive physician reports, clinical reports, etc. across multiple boundaries. Custodians of the data transmitting it, storing it, accessing it, and the transfer of patient data, but within and outside the states is just a huge management issue, which is one reason your committee is in place. Again, if you take the smart grid example, or if you also take the cloud example, you take the e-government example and a host of others, you'll see that there is the need to begin looking at privacy management holistically.

The International Security Trust Privacy Alliance in 1999, we formed. We have a Web site, ISTPA.org. We started developing some of this. The perspectives that we have are operational, so we try to stay out of the policy fights over is this policy a better policy, or is the EU model for data protection superior or inferior to the U.S. model. These are things that we felt are not resolvable at the operational or technical level. In fact, we feel that the technical side needs to be able to manage these. And we see that in the corporate community where we're, as companies, expected to manage against policies in Canada, France, Germany, Australia, and others where we have facilities with respect to PI associated with employees and with customers, so we need to have an architectural focus and technical capability to manage all of these, at times overlapping, and maybe at times inconsistent policies.

This model will be usable for system designers. It might also be usable for the informed policymaker who would like to be able to understand the design implications, the architectural implications of privacy.

We did a study, analysis of privacy principles derived from 12 international instruments. You can download this Webinar, and we developed those composite principles that I mentioned earlier just as a reference to test this model, and then we introduced the first framework in 2002 actually at Carnegie Mellon. I know Latanya Sweeney is one of the members of the committee. Since then, we have modified it significantly and reintroduced it last year as the reference model, and it really converts privacy.

Its job is to provide a structure in order to manage the full lifecycle of PI and to make the right architectural decisions. And it does that in a way that says it's not just the focus on a transfer of data from individual to a system. It has to look at the management of PII and personal health information, in your case, across multiple instances, multiple systems, multiple custodians, multiple processing, and multiple storage, and it has to have that perspective. So we needed to make this model PI and policy centric.

That is, you need to look at the personal information itself and then the policies associated with it. You need to manage multiple policy instances because you're going to be, at times, in conflict, and you have to resolve those to the extent possible through technical means, as well as business processes. And you need to look at PI and policies as objects that can be managed by systems. And the more you can manage this in a systemic way using technical mechanisms, the greater assurance you're going to have of conformity, quality, and less exception processing, which would be prone to more errors, etc.

I think it also facilitates auditing. And you need to manage this in a network lifecycle context. It isn't just one collector and one patient, now one individual with devices in a home feeding data to one energy provider. You're now looking at a network lifecycle where the data is flowing in many ways out of the knowledge and control of the patient or of the individual. And you need to integrate security services into this, whether it's identity lifecycle, federated identity, access controls, resource protection, audit and encryption, etc.

In our view, only a standards based structured model gives us the ability to implement, manage, and insure compliance with privacy and security policies, and expectations of citizens in existing and emerging infrastructures. I think the trust that citizens and patients and consumers expect can only be addressed if you can demonstrate management and demonstrate compliance. That's the overview of the reference model. If we are successful, any complex infrastructure, the smart grid, health IT, combined and supported by security infrastructure and the privacy infrastructure integrated together then using standards and mechanisms appropriate to components of those infrastructures will be successful.

Let me turn it over to Michael Willett, who will talk about the now actually talk about the model.

Michael Willett – OASIS
Okay. I assume I'm on and everyone can hear me.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Yes.

Michael Willett – OASIS
Just briefly, my biography there that was flashed up earlier indicates that my background is largely in security and security research and cryptography research. I worked a number of years in a university faculty and then onto IBM cryptography competency center, and several consulting operations in security, but all along, all that time that I've realized, as John indicated that security was a well-established discipline and all the controls were defined that, in security, it's all a matter of implementation. But I realize, at the same time, that privacy was years behind in terms of that sort of maturity that security already enjoyed. So even though security was an essential element of privacy, the rest of the components of privacy were missing.

I've worked with John all these many years in the ISTPA working, trying to transfer the methodology we use in security, in developing those controls, to the privacy side of the ledger. What I want to talk about then is how to move toward an operational view of privacy management. Now if you look at the examples that John gave, he gave examples of privacy requirements, the best practices, fair information practices, legislative instruments. All of those lists, if you will, of items are all in what we would call the domain of privacy requirements because if I handed any one of those lists to a system designer, a person that you asked to design an automated system in a computer to manage privacy based on those principles, the system designer wouldn't know where to begin.

The first thing a system designer would notice is that none of the principles mention any control. None of the principles mention any interfacing between. None of the detailed functionality is described in those principles. That's the point where we picked it up is even though the world is … in privacy principles, practices, legislative instruments and so on, where's the model, the reference model that's needed to bring that to a system design perspective.

In fact, along the way, as John indicated, one of the things we did just to practice was we consolidated the 12 instruments, the 12 international instruments that had a variety of language, and we created a composite set of privacy requirements, if you will, that John listed those 14 requirements, and we used that set … requirements, not as a complete list necessary, but as a good test of our reference model. What we found was that the reference model was able to implement those privacy requirements in a composite list, as well as a number of the more specific legislative instruments and fair information practices and so on. That development of that reference model is what I want to talk about.

There have been some privacy standardization efforts underway around the world, and John alluded to a few of those. The one that got a lot of fanfare a few years ago was the W3C work on P3P, which is called platform for privacy preferences. Basically it's a grammar for expressing privacy preferences. It's a way for two points to say, in a language that can be automated, that this is my privacy … and another point; understand these are the preferences with regard to using my personal information. That grammar was one aspect of the whole structure that's been, in this sense, standardized.

There have been some workshop work by CNISSS on privacy audit tools. There is a set of draft standards being developed in ISO now in the 2.9.1xx series. That's under workgroup 5 of SC-27, which is the security subcommittee of JCC-1 under ISO, and those standards deal basically with defining privacy requirements and best practices for privacy principles. Then there's an example, in your domain,

of the OASIS work, the XSPA, which is a standard that the federal committee developed, a standard that deals with consent directives and so on, policy exchanges between healthcare organizations. So there has been some standardization work. Once you see the whole picture, you'll realize that it's merely bits and pieces of a larger missing framework, a larger missing reference model. What else is missing? A lot is missing to do total and complete privacy management from our point of view.

First let's start with an operational view of the definition of information privacy because it's from this definition that we've started some years ago and from which we derived really all of the operational services that I'll describe. Information privacy is the assured, proper, and consistent collection, processing, communication, use, and disposition of personal information throughout its lifecycle. This is, of course, consistent with data protection principles, whatever the local domain policy requirements are, and particularly the preference of the individual.

The words proper and consistent apply throughout the PI lifecycle. It applies to all actors, all systems, and all networks that touch the information, so I'll make reference, as we go to touch points, these points in this, as John described it, the asynchronous lifecycle network flow of personal information, all those many points where the PI could be touched by an entity. What's needed then is an abstract model enabling full lifecycle privacy management at each of those touch points.

The quick evolution of the work in the ISTPA, as John indicated, we published a framework document in 2002. That has steadily evolved into the published document a reference model in 2009 of last year. Just to highlight a few things that we did to create that evolution: The first thing we ran into, and probably in your minds too, is the word framework is loaded. It's way too loaded.

When you talk to the policy people about a framework, there's a misconception that you're trying to sort of shove a policy down their throat, that you're trying to define the policy that applies to a particular domain and set the policy. That's not the objective. The objective is to define a model, a reference model that can implement in any policy that would take the policy in as a parameter for execution by the reference model.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
Is there a possibility to ask a couple questions?

Michael Willett – OASIS
Certainly.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
On the last slide, I just had—this is John Houston from the University of Pittsburgh Medical Center—where, under your operational view, is this idea of oversight and governance? Is that contained in one of those?

Michael Willett – OASIS
That's embedded. Once you're taking this lifecycle view of personal information in a sort of asynchronous network environment, at each touch point you have the question you're raising. What structure do I implement at each touch point to enforce the policy? So it comes out, I think, explicitly in the services that we define.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
Thank you.

Michael Willett – OASIS
As you pointed out, there's a real focus on that.

John Houston – Univ. Pittsburgh Medical Center – VP, Privacy & Info Security
Great. Thank you.

Michael Willett – OASIS
We validated the framework services that were defined, but we noticed that we really lacked a robust reference to network a synchronous lifecycle to PI. There was too much focus on subject and requester. Subject and requester is kind of the simple problem. As we all know, personal information moves way beyond the initial requester if combined and morphed and managed in an asynchronous fashion with a long lifecycle, so we wanted to move beyond the model of just subject and requester.

We … support use cases where that PI was disassociated from the data collector, and even the individual's control. It moves beyond the collector, and it moves beyond the individual's direct control. The lifecycle extends beyond the collector, and as John indicated policies, there could be multiple policy domains, and policies change in the future. How do you adapt to changing policies?

The other thing we did was to add more of a formalized syntax to the reference model, so we could better understand the service, the service relationship. The model, after all, defines a set of services that can be implemented and executed to perform privacy management. There's a service-to-service relationship. We wanted to formalize syntax for that. The final thing was that we did, as I indicated, we tested the reference model not only against a series of privacy, fair information practices and so on, but also against the sort of uber test, I guess, the composite privacy requirements derived in the analysis document.

Just to jump to it, even though it took a while for us to get here, and there were multiple variations on this collection of services, these are the ten privacy management reference model services, and we group them logically into three categories: policy services, presentation lifecycle services, and privacy assurance services. These are brief definitions given here in the privacy management reference model document that you can download now on the left side of the screen, if you want, or from the ISTPA Web site. You'll see more detailed definitions and syntax for each of these services. I'll give examples of that later.

In summary, the agreement service negotiates agreements between touch points. A touch point wants to use certain personal information under certain conditions, then the agreement service is used to negotiate that permission or denial possibly. That's agreement. Control is the policy engine. Policy is input as a parameter to the control function, the control service, and it provides execution, rigid execution of that policy. It, broadly speaking, provides data management.

Interaction, and again, look at this, start slipping your system design hat on from the point of view I'm moving now from the domain of requirements into the domain of actually building automated software at each of these touch points, the interaction service that handles the interface between services and exchange of notices and so on, and also the interaction between the system inside and the external world. There's an agent concept, an agent service, which means it's basically the persona for that touch point. If I've got an engine and a collection of services inside an automated point, I need an agent, sort of a personification of that point to represent that point internally and also to the outside world.

Usage is a service that deals with subsequent use of personal information and including such things as aggregating personal information or anonymizing, slicing and dicing personal information in various ways. Access allows an individual to review and suggest updates to personal information. Under the … I've got

certification, which validates the potential of participants in the whole infrastructure, and validates that certain processes are trusted. Audit is your well-known audit function, which interestingly enough is not ever explicitly mentioned in the requirements side, but audit is independent, verifiable, accountability of the operations of the reference model.

Validation is a service that checks the accuracy of personal information, and then, finally, enforcement handles the exception conditions. If a flag is raised by audit service, for example, then enforcement would be invoked to potentially redress, address the violation and perhaps even move into redressing for the violation. In a nutshell, these are the ten services. And this is more of a – it's not necessarily meant to be a system design picture because they don't relate in this way. But it's more of a schematic view of the reference model in the three grouping of services and the ten services underneath. And the very important point that John has made is that we're assuming that the security services are fully available to this model and it's arching underneath all ten of the services. That is, the ten services I've described, as you'll see in the reference model, as it evolves, all have access to the appropriate security functions, whether it be confidentiality, integrity, assurance, those kinds of functions in the execution of this privacy management service.

Let's look at a touch point, and this is just an arbitrary touch point that is a point in the asynchronous network infrastructure through which personal information, along with preferences and restrictions flow, so this is one touch point enabled with a stack, if you will. I'll call it the privacy management stack, an operational stack of the ten privacy management services. You see the agent persona encompassing the repository for personal information, preferences, and so on, the control, the agreement and interactive services, the interaction with the outside world as a touch point. The access service and then the usage service, and then I'm showing here what's called – we're calling it PI container, the PIC, which is the vehicle used to exchange personal information between touch points, along with these preferences and regulations and any contextual information relevant to that exchange.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Michael, this is Dixie. I have a question on this.

Michael Willett – OASIS
Sure.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Access and enforcement, at least, are in my mind security services. How do you see this? Where's the linkage between this and, well, this enforcement, for example, and security enforcement?

Michael Willett – OASIS
This is taking the example you gave. We're all familiar with the requirement of access. That is, individuals should have access to personal information held by an entity on themselves, and then the regulations all say with reasonable cost and so on.

Access from a privacy perspective would be the validating of that request that the individual is the person allowed to request access and validating the credentials for that person within the requirements of the policy requirements, granting access to all or part of the information, managing the flow of information from the individual to the touch point. So there's a series of activity underneath the concept of service access that is an interplay between privacy management and security. At some point, access is using security functions that exist. Maybe I'm using confidentiality for the exchange. I want to tell you, this is my real name, not that name. You've got my name wrong … so I may use confidentiality to exchange or

cryptography to exchange that. But the actual granting of that permission to the individual is part of privacy management.

John Sabo – OASIS
Dixie, it's John Sabo. Just a quick comment on that as well: There are few of these services where the relationship between traditional security services like access control, identity, and authentication on the security side map up to the privacy policies associated with individual access. When we looked at our requirements across all those instruments, and if you look at some of the control statements that you read, there is the first policy statement is, should anyone in this group or in this community or having these roles or having this set of permissions be given access at all. Once that policy question is established through the lifecycle, then you invoke the appropriate identity authentication and access controls to implement it. So this is an area where the two are very closely aligned, but somewhat distinct.

The first question is, should the individual have access to a record, for example? And that becomes a policy statement that, in a reference model, would be tagged in some manner, whether it's a logical tag or a policy relationship that would answer that question. Then the second question would be, which is more security. Does this individual have the right to see or modify or change that data?

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Okay. Yes. Clearly access, the privacy policy is only a small part of the decision rules regarding access, so this would be just kind of a handoff of privacy information to the security access control rules.

John Sabo – OASIS
Exactly, and that's a key – that's an important thing when Michael talked about integration of security into each of these services, as well as the overarching infrastructure. That's a very … key point.

Michael Willett – OASIS
And it's not so much a handoff, as you'll see in a second here. In the syntax, we've actually integrated security as a service into the fabric itself, so the security is considered to be hand-in-hand with these various services.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
This is Walter. I have just a couple of quick questions on this particular slide, and on the previous one. The first one is about the validation and when you refer to validation as the checking of the checking of the accuracy of the personal information. That's a different perspective in terms of validation in the privacy context than validation of the privacy preferences around the ability for someone to access personal information. I was a little confused about that part of the validation. That the validation in this model focuses apparently, as you describe it, on checking the accuracy of the personal information, which is a little different, I guess, then the validation of the privacy preference that controls access to the information itself. Can you comment on that?

John Sabo – OASIS
Michael, can I make a comment on that? The service definition is broader than accuracy. We're giving shorthand in this talk. If you look at the actual document, validation in the lifecycle may be something like quality of the data in terms of its relevance for a use, or it could be the timeliness quality in the sense of its immediacy. Let's say you have – again, I'm not trying to be an expert in the health space, but let's say you've got six prescriptions that have been written for a particular patient over the course of two years, and the timeliness in the case of validation would be something in that system is validating that. When you're accessing the most current prescription, you're seeing the most current prescription and not something that is two years out of date, even though it may pop up on the screen or be introduced into

the system. The validation service has many functions associated with quality, accuracy, completeness, relevance, and timeliness, which of course would be driven by policy.

Michael Willett – OASIS
Yes, I did say personal information, but as John indicated, there's a whole set of parameters that can be checked by validation.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Does that include disambiguation of the rules?

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
That was going to be my follow-up because really management requires cross-reference verification and resolution of conflict of preferences. And, in the validation, that's what I thought would be the function, so would that be included?

John Sabo – OASIS
Yes, because the function – don't forget the function may be – I had mentioned this earlier in my part of the talk. To some degree, it's desirable to automate as much as possible. To some degree, conflict resolution may not be such that it can be automated and will require human intervention, which was part of a business process. That is also accommodated in this service.

In several of these services, you will have to have disambiguation because, for example, the control service. You may have conflicting policies. You may have two entities that in fact are accredited and generally meet requirements for assuring PII or health information, but there may be certain aspects of their policies that cannot be adjudicated in an automated fashion, and some additional processing or agreement language has to be built into it. I think, in several of the services, disambiguation will be required.

Michael Willett – OASIS
You'll see, in general, when I drop down to the syntax, that the reference model is really at a reference model level in terms of services. We have not designated the particular mechanism that could be exploited under each of these services, so that's work to be done. So it's still at, in a sense, the reference model level. And some of the things you're talking about are particular mechanisms that would be available to validation or to some of the other services in a particular actual implementation.

John Sabo – OASIS
I think you'll see when Michael moves to some of the later slides on next steps that some of the issues you're raising must be addressed. Again, this is a fairly abstract model, which we think is a very important one, but it also raises questions that will need to be addressed in OASIS and through use case development.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Thanks.

Michael Willett – OASIS
And notice on this chart that there are … the backdrops here are the legal, regulatory, and policy contacts. Those become input parameters to the ten services. Secondarily, the security foundation is always there, as we've indicated with access, which is pretty strongly less so with some of the other privacy management services. The security foundation is always there, assumed to be there to be called on when needed and appropriate.

Now if I just connect, if I just show the arbitrary connection of two of these touch points somewhere in this network in an asynchronous infrastructure is that each of these touch points in different regions, domains, with different priorities and requirements and policies possibly all have basically the same stack, if you will. They are parameterized or customized or configured differently depending on the context they reside in. And so the agreement work between two touch points, the exchanges of … happen when appropriate, that sort of thing. They both have access to their own set of assurance services. This is kind of the high level operational view of any two touch points in the infrastructure.

Now I want to look at a very, very, very simple use case of this in which I have left out some detail, but I will allude to here in a second. The basic application is a subject and requestor, even though our focus is really well beyond this subject/requestor paradigm. I have an employee and a business application like payroll for example, and the employee application like payroll request certain personal information from an employee. I'm assuming that I have two PMREM stacks implemented, one at the employee and one at the payroll application, configured differently obviously with different policy parameters and various other input.

The basic flow here, the basic flow would be that the agent through interaction on the payroll side would present a notice of the purpose used and requested PI to the employee on the employee side. I'm leaving out perhaps negotiation arbitration here. …permissible purpose and use would be stored in the database on the employee side by control, and then transferred to payroll where the PI is probably subjected or validation of what other independent capabilities the payroll application may have to validate the personal information and then store it in the PI database on the payroll side under control, so just a basic interaction and flow of service interactions between the two stacks.

Of course, even in a simple application like this, I could well have invoked audit on both sides. Every action performed on personal information by a control function should be audited, should be logged, and I may have configured audit with certain not permissible uses pre-configured with not permissible actions. Those would have raised triggers or flags to the enforcement function. I could have asked that the credentials of payroll be certified before I do any dialog with payroll. There are a number of richer interactions with services that could be invoked, even in this simple example. I just wanted to give an example of how this all works.

Then the other thing is on payroll, the personal information, even with the permissions granted, could well move beyond the database and the storage of payroll in the subsequent applications, either by explicit or implied permission of the employee themselves. There is, again, the implied infrastructure distribution of personal information, even in this picture.

I go back to one of John's network, asynchronous data flow pictures, and just point out that what we're after is, in ISTPA and OASIS, is a reference model. We haven't designated the mechanisms yet and the specifics under each service in terms of mechanisms, but a reference model that can be built and configured at each of those touch points, whether it's requestors, users, applications, aggregators, and so on. Each of those touch points and beyond in that asynchronous infrastructure could be implemented, configured, and personalized and, therefore, handle the exchange of personal information and preferences throughout the infrastructure.

Let me drop down then now into a little of the syntax that we've developed so that the reference model is a little more structured and also, we found, amendable to modeling and simulation. There are ten services. For each service, we've gathered a collection of functions. There's a function set for each of

the ten services. The function sets, there are seven function sets for each service. Now they're different for each service, but categorically they're the same seven function sets.

A defined operational requirement select the parameters, input the necessary parameters, process the data and parameters, output selected parameters and values, link to other services to complete the operation, and then, underneath all this, secure that operation of that service with the appropriate security function. These are seven function categories that we found very complete, in fact, and useful in defining each of the services.

A use case, for example, would invoke a sequence of service calls. Think like a system designer now. One service … the exchange of a notice and so on. So services are calling each other in a call convention. Then once a particular service is called, that service will execute a sequence of these seven functions to execute the use case. Let me just go into two examples. I pulled out from the ten services two of our more relevant services from the ten services agreement.

Here's the definition of agreement from the document itself. Basically if you'll glance through the agreement, it deals with, as I have indicated earlier, negotiating between touch points, the result of which can be consent, denial, or a derived agreement among the parties. This is between any two touch points, not just requestor and subject.

Control, another service function is that control engine I mentioned that governs and executes the policy. Managed in accordance with prescribed privacy policies and controls. That's the basic control function definition.

Let's look at each of these internally. If I look at the seven functions under the agreement function, for example, there's a common define the agreement requirements, so like … parameters. According to the way we kind of detailed out the definition of agreement, there were two input functions defined, so there could be multiple instances of these function categories depending on the service. There was agreement of PI definition and agreement for additional permission input.

There were three process functions, and you can think of processing as kind of the guts of the service itself. It's where things actually get done. There's input, output, and security, and then there's the guts of the service, which is process itself. There were three process functions defined for agreement. I just detailed out, these are all, by the way, defined at a lower level down.

And I just detailed out one of those process functions, process agreement exchange. That's the exchange of initial parameters related to a potential agreement between parties. The idea here is with agreement, I can make initial proposal. Here's a conditional proposal for you to exchange information with me, and then I may call the process agreement interchange, which does subsequent exchange or bartering, if you will, to reach a final derived agreement or a rejection of the proposal. That green area on the left shows you just the next level down of definition for these process functions.

Then, of course, the agreement continues on through the output, either output consent or denial or an agreement among the parties. Linkage, the linkage function with link agreement to other services, as needed. And then security then would invoke security controls in support of the agreement function as appropriate. That's sort of the constitution, if you will, of the agreement function. They all are very similar. They just differ in the details of those seven function categories.

For example, if I look at control, the same seven function categories. I had the define, the select, the input. There are two inputs here, two input categories, if you will, so two input functions are defined for

control. There are, again, three process functions: control configuration, control management, and control of PI input and output under control, and each of those, again, functions are defined at another level down in the reference model document.

Then, from there, I move into the output. There are two output functions for control. There's the linkage to other services and then the under-arching security control functions available to the service. That's just an example of two of the ten services and the PMRN syntax section goes into the complete detail of defining this syntax for each of the ten services.

This is one way of viewing sort of the cycle of going from a set of privacy requirements to an operational implementation of privacy management control. I start at that little yellow picture, the little yellow square. I have a set of privacy principles and practices. Perhaps combined with privacy laws and policies. That gives me the requirements input, right, the requirements input to this operational aspect where the privacy management and reference model takes hold. I select the reference model services and functions that are particular to that use case in particular. Then I select the security functions that are needed in implementing that use case.

That then moves out of the reference model domain and into a lower level derivation of what the system design people call an architecture. What I've described here is a reference model. It's not really an architecture in the detailed sense. And underneath the architecture would then derive the operational controls themselves, the actual privacy management mechanisms and controls.

The other way to look at this, by the way, is more of a linear view. Let's look at this chart as I move through four stages of developing a privacy management implementation. I have a defining stage where I've got the input from the requirements and practices and laws and that sort of thing. That moves into the operational definition stage where the privacy management reference model is invoked.

I select services. I select the security requirements and services. Below that, I move into developing an architecture for privacy and security architecture, combine security for a privacy and management security architecture, and ultimately then into a low level mechanism. Do I use SSL here for confidentiality? What sort of database access controls do I invoke? What are my identity management, identity and trust management functions that I invoke here? What are available? What can I implement? Those are all the mechanisms that are selected ultimately to build the privacy management solution.

The next steps, John alluded to some of this. I mean, that in a nutshell then is sort of a quick overview of the reference model. You need to see the reference model, in our view, as bridging the gap between the requirement side of the ledger and the implementation side of the ledger.

We haven't gone as far, as I say, as implementing the particular, selecting the particular mechanism or even a very rigid architecture for this. That's still flexible. But the important step, we think, was going from requirements to an operational view, to a system design reference model view.

Our next step in the OASIS process are to create the technical committee, the proxy management reference model technical committee. We've been soliciting initial proposers of this, and we've got a number already onboard that are behind the reference model, including NIST, their interest in the smart grid, for example. The American Bar Association, other organizations and individuals have already signed on, and we're soliciting others.

We've drafted a technical committee charter already. We're planning the first technical committee meeting probably in Washington, D.C. We're looking for partners now, in the interim, and also ongoing to

host workshops to test the reference model against various use cases, and that's what we've been doing in ISTPA, and we'd like to have more help in that regards to take vertical disciplines and test the reference model in those various disciplines. That's basically the objective of the TC is to refine the privacy management reference model and develop and solicit use cases to test the reference model, ultimately moving, we hope toward an industry standard at the reference model level.

We worked with the smart grid people, the OASIS blue member section and NIST, and most importantly, as John indicated, is being able to work with you and other members of the health IT community on relevant use cases. With that, I think we're moving to a period of questions. If you have questions now, or if you want to send us questions or comments or whatever, here are the e-mails for all three of us from the OASIS group. With that, I thank you for the attention.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Thank you, John and Michael. We've asked a few questions as you've gone along, but are there others that people on the line might have?

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Yes. This is Walter Suarez from Kaiser Permanente again. I do have a couple of questions. First of all, thank you so much for this overview. I think you have taken us a few steps higher, above and beyond healthcare into other industries, and even beyond U.S. boundaries into international territory, looking at perspectives and models from other countries and building what you are building as a reference model.

I think one question I have is around the concept of, in order for all these to work, policies, and when you say when you use the word policies, some in the policy arena could read laws and regulations whereas someone in the IT arena would read beyond that and look at system policies, as well as organizational policies, not just regulations and laws. But above all, basically policies, in order to be able to be managed by systems, I think one of the, as I read it, conditional elements in your model, they need to be seen as objects. And in order for them to be seen as objects, they need to be codified and be able to be computable. Is that a fair expectation that in order for this framework to work appropriately on the input side, on the policy side, it needs the policies to be computable and codified?

Michael Willett – OASIS
That would be the ultimate holy grail, in a sense. And there's work in that direction. P3P, for example, is an automated language for formalized language for expressing privacy preferences, and it's no necessarily all encompassing yet. It's expandable.

There are other works in that area. XACML and others are writing languages, giving us formal languages for expressing policy, so that's the ulitmate holy grail is that the policy itself as a parameter can be fully expressed in some automation language. But if not, as John indicated, this is a reference model, and whenever the reference model can invoke human interaction, obviously, and other procedures that aren't as automated when those are lacking.

John Sabo – OASIS
I think you can say something like this, that one particular value of the model that we have seen in discussing it with practitioners, often people in the security community who don't really understand data privacy or that distinction is that it opens thinking in terms of it opens your dialog about what policies are needed across these boundaries. By policy, it could start at a fairly high level of a set of policy requirements. It could be, for example, if you looked at validation, the validation service related to data quality. There may already be expectations of data quality at each stage of a health IT system. The question is, in a lifecycle view, you know, model view across all these boundaries, do you have in place

policies that would allow you to manage that level of quality consistent across all these boundaries associated with the personal health information.

One value of this model is it forces the thinking in the architectural community about managing and defining these policies at a fairly defined level, not totally abstract, but across all these boundaries. That's level one. The second would be what mechanisms then can be architected to manage those policies. And then, as Michael said, to the extent that they can be computable or made automatable, can that be done because we have expression languages to manage it. In some cases yes, in some cases no, they would have to be enforced through contract or agreement language on some kind of monitoring, you know, auditing to validate that the system is working as planned.

I think the model doesn't mandate technical instrumentality for every aspect of it, but the reason we and, Walter, your point about the higher level looking internationally, usually, I mean, often enterprise systems and architecture investments may or may not build totally separate infrastructures, particularly when you look at the security aspects of it or data storage aspects. You may have to restrict controls about how data is stored and how it's accessed. But the infrastructures may be an enterprise infrastructure. For that reason, we felt it was really important to design a model that would accommodate any kind of architectural design, accommodate different policy boundaries, accommodate different requirements and not be inflexible.

I hope that explains a little bit. It does not require computable, but that is kind of the holy grail. It's likely not going to happen, and in some of the services, the boundaries between physical engagement, particularly in the enforcement services as an example, and what can be automated, that boundary is kind of permeable at the moment. And that's one reason we need to do more work, and that's one reason this OASIS work is very important to address the kind of question that you ask. We need to tease out some of these issues and determine if the model needs to be adjusted or how do we accommodate these things in a certain use case.

Michael Willett – OASIS
Yes. We've already done a lot of adjustment with the model based on our own internally derived use cases that we've played with in the ISTPA, so that's why we're really soliciting real life use cases because it'll help tease out more of these nuances.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Yes. Just as a follow-up, I think it would be very valuable to evaluate a healthcare specific use case. As I said, this has been great to bring this perspective of beyond healthcare, you know, all these other industries that you're looking at from a personal information or personal identifiable information perspective rather than an identifiable health information perspective. It's great to see that.

Then it's also great to see the international perspective, but as we say in healthcare, all healthcare is not even anymore local. It's personal. And nothing more personal than privacy, of course, and what I wanted to bring in is, you know, it gives me the impression that the model is attempting to be architectural, operationally, and technically agnostic to the jurisdiction or the jurisdictional environment in which it operates. In other words, it's capable of handling any and every possible policy environment.

As you might be aware, of course, in healthcare, in this country, and probably in other countries, there are many laws and regulations that control many aspects of health information, and those laws are not just national or federal in the context of the United States, but also state and local regulations. And so, being able to be agnostic to accommodate multiple privacy regulations and ultimately privacy policies is another

holy grail, I suppose, and some people believe that it's almost impossible to do that without this ability to codify and to make policies objects that can be computable.

John Sabo – OASIS
The model is agnostic for precisely that reason or attempts to be as agnostic as possible for precisely that reason. But the development or the application of the model to an environment is not going to be agnostic. In other words, the model leads to the issues you just talked about where you've got local, state, you know, you may have policies flowing from the medical community, from HHS regulations. You may have practices. You may have existing languages that are used to convey health data to providers or to code. You have all these things in place already, which gets right back to the value. So the model itself is designed to be agnostic, and we ran into some issues early on with that with some – you know, from policymakers who insisted that it be agnostic because they felt that it should not have a bias, but the usability of it. And, for example, if you, if members of the committee or perhaps the working group, the standards group, or practitioners in your committee and others wanted to work in the structured OASIS environment as an offshoot to do a use case that's applicable to HIE or whatever particular focus you wanted to bring to it, that's precisely what has to happen.

Then you may end up with a health IT profile, which is specific to circumstances that you're developing that would basically still fit within the model because the model itself provides these overarching services, but the particular mechanisms, for example, we've also heard feedback that you've heard of the term capability and maturity model that there's a level of implementation that is very advanced. It may be less advanced. That may be accommodated in the model as well.

It may be that some of the mechanisms are two or three years off because a standard has not been developed. I think the use case development component of what we're trying to do, and I know one of the things we were talking about in the technical committee is reaching out to organizations to develop specific use cases and work with the model and then help define it for particular communities of interest or particular infrastructures. That's precisely what we have to do.

Michael Willett – OASIS
Yes. To say that privacy, I mean, to say rather that policy is a parameter, I mean, it's easy to say. It's an oversimplification of what actually has to transpire, as John indicated. In each jurisdiction, there has to be a way to express or personify the policy to the reference model itself, whether that's manual or control or automated, whatever. There's a lot of meat missing in the mechanism level. But, in the agnostic sense, we say policy is a parameter.

John Sabo – OASIS
Now you may decide in your use case, for example, that – I mean, you might look at this so incredibly complex environment, which deals with the patient trust and then multiple providers and so on, and payment systems, and HIEs. You may say, well, in the model we need to develop, we've evaluated what's available that are mechanisms to support certain services. We're looking at this in a very holistic lifecycle manner. But we've determined that certain functionality will have to be addressed through work later, and then that basically says you're using the model to help design the architecture, but the architecture itself has to accept certain limits based on where you are in the development of standards or development of practices. Yet the roadmap can be there from a capability maturity model perspective to advance this additional work in the context of this lifecycle reference model.

Honestly, one of the things we've seen is that it's been an eye opener for so many people who typically view privacy as, well, I provide – in my case, I provide, my GP gets my data. It goes into a little folder or maybe go into a PC. I have no idea what happens to it. I just had a test with a specialist. I have no idea

if the data flowed back to my GP. I just, as a patient, I have no transparency about any of this process. A reference model at least gives us the ability to begin building trust and transparency.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Yes. Thank you. I yield back to our chair.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Sort of related to Walter's question, which was a real good one, a lot of our effort is focused on exchanges, and it seems to me, in any of the domains that you might apply this, the idea of being able to share attributes and policies between organizations, between jurisdictions, being able to even share the context of the permissions is important. And I don't see any of these services addressing that.

John Sabo – OASIS
Two services – I'll let Michael get into detail if we need it, but the control service, the usage service do address that. In other words, to use your example, if there is a rule related, if there are permissions associated with particular personal health information, and you bind the personal health information to that permission in terms of a rule, and the rule flows, then the control service is the initiator of that. And when the data, the PHI or whatever moves to another entity through another exchange or through a participant in the exchange, the rule is still associated with that, but it's managed by the usage service because the usage service deals with data management beyond the originator of that control function, if you know what I'm saying.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Yes.

John Sabo – OASIS
This is an area though where the use case development will help inform how the services actually operate and what mechanisms. I know Michael referenced the so-called personal information container. It's just an abstract concept. It could be a logical link. It could be some kind of digitally signed object. It could operate in any number of ways, but the point of it is that it's managed by control and usage over time where it's implicated in those services.

Michael Willett – OASIS
Also implement, you could use agreement between those exchanges also to sort of negotiate, in a sense, meta-policy between two exchanges.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Okay. Yes. Agreement, yes, the agreement need not be at an individual level. Yes.

Michael Willett – OASIS
No, that's right. Exactly.

John Sabo – OASIS
No, that's right. In fact, we had pushback on agreement because originally we had called this the negotiation service, and we took a lot of grief because, in some instances, frankly, there is no negotiation at the individual level, and people were looking at it only at that level, so we looked at the end state, which is an agreement, and it could be entity-to-entity agreements, system-to-system agreements.

Michael Willett – OASIS
Or denial.

John Sabo – OASIS
Or denials or a patient to a provider.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Okay. Are there other questions, comments?

Michael Willett – OASIS
I would encourage those that have time on their hands to download the two documents from the ISTPA. I think you'll find there's some … materials for motivation in the front of the PMRM document, not too long, and the analysis document is a real eye opener in terms of the disparity around the world.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Yes. And I would point out once again that the presentation itself, the reference model, and that ISTPA analysis of privacy principles all can be downloaded on the left side of your screen. Deven, I think, in particular, the analysis would be useful for the policy workgroup.

Deven McGraw - Center for Democracy & Technology – Director
Yes. No, I totally agree. I'm actually looking for that left side of screen link that you're talking about.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Did you blow up the slides maybe? Yes. They just moved it.

Deven McGraw - Center for Democracy & Technology – Director
Got it. Okay. Yes. Thank you. Thanks for the cue.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Yes. Thank you.

John Sabo – OASIS
Dixie, one thing I would also say is that we expect to initiate the new technical committee within OASIS in May, if all goes well, or end of April. And we would very much like to organize subcommittees to look at use cases. We're already discussing with the smart grid, OASIS blue member section to do something like that. If ONC or the committee or members of the committee or other entities who are building these systems want to engage to actually create a workshop or a set of workshops or a subcommittee focusing on use case development for health exchanges or some aspect of the health IT infrastructure, that's absolutely what we want to do, and we invite participation in that.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
We will talk a little bit more about that.

John Moehrke – Interoperability & Security, GE – Principal Engineer
Dixie, this is John Moehrke. I have already signed onto this as an individual member, so I do invite others to come in as well because certainly one individual is a very difficult thing to go through all of the analysis. I think this is a nice framework to examine the currently healthcare work in this light. I know we've had a similar discussion in the HL-7 security working group where there is work to capture the computable form of a privacy policy in a CDA document, which of course would be that missing piece that Walter uncovered here. We are working within HL-7 on a way to encode the privacy policy so that we can move it around, and we need to make sure that we run this particular model against that to see if there are any other gaps that we need to fill.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
That's great, John. I was going to actually ask the question of how does this model relate to the HL-7 work being done by the security groups, so thank you for those comments.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Actually, that's the third of our four presentations is the HL-7 working group. But I did think, and I think your presentation has confirmed it that this was a really good place for us to start.

Walter Suarez – Institute HIPAA/HIT Education & Research – Pres. & CEO
Excellent, yes.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Because it really lays, you know, a nice framework for everything we'll be discussing with the other standards with respect to the other standards work. If there are no other questions, let me thank John and Michael and Dee. This is an excellent presentation, and we really appreciate your working with us on this.

Judy Sparrow – Office of the National Coordinator – Executive Director
Dixie, we might want to open it up to the public to see if we have questions on….

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Yes, I forgot about that.

Judy Sparrow – Office of the National Coordinator – Executive Director
Yes. Operator….

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
As I mentioned to you and John, the last 15 minutes are so, 10 to 15 minutes, we open up to the public. The public hasn't been able to speak thus far, so that's what we're doing now.

Judy Sparrow – Office of the National Coordinator – Executive Director
Great.

Operator
We have a public comment from Phyllis Johnson.

Phyllis Johnson – New York
Hello. This is Phyllis Johnson in New York. Just a technical question: Where are the Webinar presentations available, because when I tried to log in to the Webinar, it said the meeting had been adjourned?

Alison Gary – Altarum Institute – Communication Technologies Coordinator
The materials for this meeting are at the URL altarum.acrobat.com/ps. They're also available on the HIT ONC FACA Web site.

Phyllis Johnson – New York
Thank you.

Judy Sparrow – Office of the National Coordinator – Executive Director
Are there any other comments from the public?

Operator
We have one more comment from Ashok Malhotra.

Ashok Malhotra – Oracle – Consulting Member Technical Staff
Yes, this is Ashok Malhotra from Oracle, and I have a comment about your reference model. This is supposed to be … model that covers lots of privacy situations. The problem is, it's a very strong model. It's a very strong model because when you actually request it, you have to … why you want it, what you're going to do with it. It is a starting interaction between the parties, and after that, you can actually understand the data. Often it is possible to actually do a simpler model where, when you actually send the data, you send policies with it. So I'm just wondering if your reference model handles this simpler case.

Michael Willett – OASIS
Again, what you're describing, it sounds to me like whether you can either push or pull policy.

Ashok Malhotra – Oracle – Consulting Member Technical Staff
Yes.

Michael Willett – OASIS
That's possible. The example I gave was notice was presented to the subject, but you could implement the services and configure them differently in your use case or your domain. The policies are drawn, are pulled, if you will, and not pushed.

Ashok Malhotra – Oracle – Consulting Member Technical Staff
They're actually on with the data.

Michael Willett – OASIS
Yes. Ultimately you want the permissions to be logically associated with the personal information throughout its lifecycle.

Ashok Malhotra – Oracle – Consulting Member Technical Staff
Absolutely.

Michael Willett – OASIS
However, that's what we call logically the PIC, the PI container.

Ashok Malhotra – Oracle – Consulting Member Technical Staff
Yes.

Michael Willett – OASIS
…logical concept. You don't need the remain together, bound together tightly in some fashion.

Ashok Malhotra – Oracle – Consulting Member Technical Staff
Yes.

Michael Willett – OASIS
But whether they're pushed or pulled or whatever is really a policy question.

Ashok Malhotra – Oracle – Consulting Member Technical Staff
Okay. I wanted to make sure that this, I think, somewhat simpler mechanism was part of your reference model.

Michael Willett – OASIS
It is. In fact, the other point is, as I indicated in the example, not all the services need to be invoked in every situation.

Ashok Malhotra – Oracle – Consulting Member Technical Staff
Absolutely.

Michael Willett – OASIS
We envision the services kind of evolving. You could probably select a core set of services needed as a basic implementation and then evolve from there. Or in particular use cases, may not invoke any of the, you know, may only invoke one or two of the services depending on the jurisdictional requirement. It's flexible in that sense. You don't need the whole stack, if you will, for every use case.

Ashok Malhotra – Oracle – Consulting Member Technical Staff
Very good.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Good question. Thank you.

Operator
We do not have any more comments from the public.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
All right. Thanks to everybody who dialed in, and thanks once again to our presenters.

Michael Willett – OASIS
Thanks to your committee.

John Sabo – OASIS
Thank you very much for the opportunity to present to you.

Deven McGraw - Center for Democracy & Technology – Director
Thank you, Dixie.

Dixie Baker – Science Applications Intl. Corp. – CTO, Health & Life Sciences
Talk to you later. Bye.

Public Comment Received During the Meeting

2. What are the legal ramifications of protection of personal health information?